

Toward a New Paradigm for Airborne Surveillance Radar

Daniel R. Fuhrmann

Electronic Systems and Signals Research Laboratory
Department of Electrical Engineering
Washington University
St. Louis, Missouri 63130

Abstract

Traditional radar signal processing and space-time adaptive processing are carried out in the radar-centric coordinates of range, Doppler, and angle. We propose here a new approach for radar signal processing which is carried out in the geographical coordinates of the environment. While such an approach may present formidable computational challenges, several advantages accrue, including: 1) the geographical coordinates of the intrinsic coordinates for the surveillance quantities of interest, especially in the GMTI setting, 2) geographical coordinates are invariant, unlike the position, orientation, and velocity of the platform, 3) using fixed parameter coordinates facilitates the integration of multiple measurements and multiple look angles, and 4) using fixed parameter coordinates facilitates multiple sensor fusion, since all sensors would be collecting information about the same quantities of interest. We describe here a research problem utilizing this approach, which incorporate elements of both space-time adaptive processing (STAP) radar simulation using terrain elevation data, and an active-testing method for surveillance radars with agility on transmit.

1. Introduction

This paper documents partial results from a feasibility study, carried out in the Department of Electrical Engineering at Washington University in St. Louis, on advanced methods for radar signal processing which make use of important side information such as that available from digital terrain elevation maps and other geographical information systems.

The initial focus of the project on was on leveraging existing adaptive detection and structured covariance estimation algorithms for space-time adaptive processing (STAP) [1,2] detection of ground moving targets (GMT), and to develop new algorithms where geographical side information is available. This required that we gain experience with acquisition and manipulation of such geographical information and to develop simple models for radar signals which are built on such data. What quickly

became apparent in this exercise is that the idea of a stationary clutter covariance matrix, for which there exist multiple independent samples, does not really fit the scenario at hand. Much more appropriate is the idea of a clutter reflectivity in each resolution cell on the ground (such as in radar imaging), which can be estimated either online or offline, and which can be used to model the interference which is competing against a target also set in the ground coordinates. Arguments for and against radar signal processing in ground coordinates are presented in Section 2 below, and results of our radar simulation work are presented in Section 3.

A secondary focus of this research effort was on algorithms and methods for exploiting the electronic agility of modern radars to adapt their transmit patterns to maximize the system's effectiveness. As we gained experience with knowledge-aided surveillance through the use of radar simulations and geographical data, it became clear to us that this secondary "adaptive-on-transmit" problem was fundamentally more interesting and more important, and accordingly our interest shifted to what we now refer to as *active-testing surveillance systems*. Here we postulate that the goal of a surveillance system is to minimize the conditional entropy, or uncertainty, regarding the state of nature in the region under investigation, and accordingly one should choose measurements which maximize the mutual information between the state of nature and those measurements. For example, in the multiple target detection problem one might wish to adaptively define beampatterns which place the most transmitted energy on ground pixels that are most interesting or require the most attention. An algorithm for doing exactly this is derived here, which is applicable across a wide range of distributions for the received data. This aspect of our work is described elsewhere [6].

The results of our work, while preliminary, are both encouraging and intellectually stimulating, and suggest several lines of follow-on research. Specifically, we would like to join the two lines of research activity, those of radar simulation using geographical information, and active-testing surveillance. We believe that it is possible to place

the active-testing methodology within the context of an airborne radar employing STAP processing and GIS/GPS/INS side information. The proposed algorithms may require extensive simulation tools as an integral part of the signal processing itself, for generating hypotheses, sampling from the posterior distribution, or some other form of analysis-by-synthesis. This suggests a mission role to be played by high-performance commodity graphics engines.

2. Signal Processing in Ground Coordinates: Pro and Con

One of the main tenets of our work, both current and proposed, is that the objective of surveillance is to gather information about ground phenomena which can be precisely geolocated. The convergence of enabling technologies in global positioning systems (GPS), inertial navigation systems (INS), and geographic information systems (GIS) makes possible the precise location of the radar platform in a local (or global) coordinate system that describes the region under investigation. Since the radar will most likely make repeated measurements of the region from a variety of positions or look angles, it makes sense for us to consider using the invariant coordinates of the ground as the coordinate system for the processing. This is in contrast to what is done in traditional radar signal processing, where variables of interest are described in the radar coordinates of range, Doppler, and angle. We discuss here some of the advantages and disadvantages of such an approach.

The most compelling argument for using ground coordinates, it seems to us, is that they are the intrinsic coordinates for the quantities under investigation, and are invariant to changes in the radar platform. For example, we might tile the ground with rectangular patches defined by a GIS or terrain map, and the variables of interest might comprise a binary target vector, where a 0 or 1 in patch i would indicate the absence or presence, respectively, of a moving target at that location. Relating the raw data to these patches seems to us a more natural approach than first processing the data in radar coordinates then applying some transformation to the results to convert the results to ground coordinates.

A second reason in favor of ground coordinates is that they are not only intrinsic, they are invariant, meaning they are fixed for the entire data collection interval. This is in contrast to the radar, whose position, orientation, and velocity constantly change. In order to integrate multiple observations into a single inference on the state of nature θ , it seems prudent to us to maintain a parameterization of θ which does not change from measurement to measurement.

The concept of using multiple measurements or look angles to perform inference on some object is of course not new. It is at the core of algorithms in medical imaging, such as computer-aided tomography, in which X-ray attenuations are reconstructed from multiple projections. In fact, the connection between computerized tomography and radar imaging was made explicit in a landmark paper by Munson *et al.*[3]. More recently, we have developed methods for generalizing the well-known MUSIC and MVDR algorithms from sensor array processing to the multiple look-angle case [4,5]. Again, the key to making this work is to use a coordinate system attached to the observed scene rather than the moving observer.

Of course, arguments can be made *against* using ground-based coordinates. The most obvious to us is that it runs counter to 50 years of radar signal processing history and accumulated experience. The very possibility of a ground-based approach has only recently become possible with advances in technology for precise geolocation of the platform combined with descriptions of the local geography, and thus it is no surprise that, to the radar engineer, the "natural" coordinates would be range, Doppler, and angle. Development of an entirely new paradigm for radar signal processing could be costly and time-consuming.

A second advantage of using radar coordinates comes from the computational advantages that accrue from the regular mathematical structure they afford. The Fast Fourier Transform (FFT) algorithm is a staple of modern digital signal processing, which can be used quite effectively in both pulse compression and Doppler filtering. Irregular terrain features, will quite likely require that we make a clean break with the FFT - a break which may be quite difficult for some to accept.

The use of ground-based coordinates requires accurate and current data about *both* the radar system and the local geography. That is, not only will one need to calibrate (or autocalibrate) the radar system, but some method for refining or verifying the position of the radar within the local geography, as well as refining or correcting errors in the terrain data, will be required as well. Here we envision a system of continual or regular autocalibration in which the data support a consensus view of the the platform position and orientation, the radar array manifold, and the local terrain. This seems to us a challenging problem, but not an insurmountable one.

Finally, the use of ground-based coordinates may require extensive rethinking of the high-performance embedded computing implementations of proposed algorithms. Current implementation efforts may focus on radar signal processing algorithms in which covariance estimation and sample matrix inversion, with indices derived from radar coordinates, are the key steps to be

pipelined and otherwise mapped onto multiple processors. It is not clear at this point whether such structures will be appropriate for the kind of ground-based processing we envision, which may involve graphics operations such as coordinate transformations, ray tracing, hidden surface calculations, and the like. We see a possible role here for the kind of technology one finds in flight simulators, virtual-reality video games, and other applications of commodity graphics engines.

3. Radar Simulation Using Geographical Data

A substantial effort was put into the development of radar simulation tools which make use of publicly-available geographical information. It was understood that such an effort would be duplicating those of several other excellent laboratories; nevertheless it seemed to us worthwhile as an exercise in educating ourselves about geographical information systems and the signal processing issues involved in modeling the radar in ground coordinates.

A preliminary investigation was carried out into the various types of geographical information available over the Internet. The two primary sources are the U.S. Geological Survey (USGS) and the U.S. National Imagery and Mapping Agency (NIMA). The USGS provides elevation data in three different standards, and land use data in two different standards, for the entire U.S. NIMA has elevation data for the entire world in the form of Digital Terrain Elevation Data (DTED) maps. For ease of use and manipulation, we chose to work with National Elevation Dataset (NED) and National Land Cover Dataset (NLCD) data from the USGS.

All data were obtained from the U.S. Geological Survey (USGS) Seamless Data Server. Files conform to the ArcGrid file format, for which one can find freely available application programs for interpreting the data. Data points are specified in a geographic projection based on the NAD83 datum, the GRS80 spheroid earth model, and the NAVD88 vertical datum. Using this 3-D earth model, each data set was recast into a local coordinate system for use in our simulation, and the data was imported into MATLAB. For demonstration purposes, three geographically interesting datasets were chosen: the Grand Canyon, San Francisco Bay, and an unnamed mountain.

In the radar simulation, a radar platform is specified in terms of its position, orientation, and velocity relative to the earth's surface in our local coordinate system. Using the terrain elevation data, the terrain surface is broken into a set of rectangular patches of size commensurate with the radar resolution (30 meters). A patch here is defined as the region formed by four adjacent elevation data points. Each patch is determined to be either visible or hidden

using a hidden surface removal algorithm sometimes known as the Z-Buffer Algorithm in rendering and simulation tools. Patches are re-projected onto a coordinate system center at the platform with the z -axis located on a line-of-sight. If one patch has a larger z -distance to another, and the center of the more distant patch is within a threshold distance of the closer, then the further patch is hidden and does not enter into the calculation of radar returns.

For each visible patch, the range, the vector velocity relative to the platform, and the area projected perpendicular to the line-of-sight are computed. The simulated radar pulse is an omnidirectional sequence of linear FM chirps of the form

$$s(t) = A(t) \cos(2\pi f_c t + \pi \alpha t^2) \quad (3.1)$$

where f_c is the carrier frequency, in Hz, α is the chirp rate, in Hz/s, and $A(t)$ is the amplitude modulation which was assumed to be either ON or OFF, i.e., $A(t) = 1$ during the pulse transmission interval T and is 0 otherwise. Specific parameters are:

$$f_c = 10\text{GHz} \quad (3.2a)$$

$$\alpha = 1\text{MHz}/\mu\text{s} \quad (3.2b)$$

$$T = 10\mu\text{s} \quad (3.2c)$$

In this model the pulse bandwidth B is approximately 10 MHz (range resolution 30m, see above) and the time-bandwidth product is thus $BT = 100$. We transmit a sequence 16 such pulses in one coherent processing interval (CPI), noting the change in distance to each patch for each pulse within a CPI due to the relative velocity between the platform and the patch. For simplicity, the angle between the platform and the patch is assumed to remain constant during one CPI.

The simulated return for a single patch is based on a random reflectivity model wherein the complex amplitude of the incident pulse sequence is multiplied by a circular complex Gaussian random variable with standard deviation proportional by the reflectance of that patch, as seen by the radar platform. We use a simple Lambertian reflectance model, wherein the reflectance is directly proportional to the area of the patch projected onto the radar line-of-sight. The land use as specified in the NLCD data also played a role here, although our assignment of reflectivity to land use was quite arbitrary and not based any known results of careful scientific study. To account for relative velocity, the return for pulse n within a sequence is multiplied by the linear phase shift term $e^{j\phi}$ where

$$\phi = \frac{2\pi f_c n \delta D}{c} \quad (3.3)$$

δD is change in distance from one pulse to the next, and c is the speed of light. The vector of such phase shifts, indexed by n , forms a 16×1 Doppler vector associated with each patch.

The effect of quadrature demodulation and the pulse compression is also included in the simulation. The RF return is *not* simulated, but rather the pulse sequence as seen through the pulse compression filter whose impulse response is the same as $s(t)$. The simulated return is sampled at a 10 ns sampling period (sampling rate $f_s = 100$ MHz) over a 30 μ s interval, leading to 3000 samples for each pulse. The collected returns for all 16 pulses in a CPI form a 3000×16 complex matrix. The returns for all the visible patches in the model are simply summed together into one radar received data matrix. Note that here we are a simulating a single receive antenna, although the extension to multiple antennas, leading to the usual STAP data cube, is straightforward.

A demonstration example of our radar simulation tool is shown in Figures 3.1-3.2 below. In Figure 3.1, we show terrain relief for a small area of the Grand Canyon as represented by the USGS NED terrain data and displayed using MATLAB surface graphing tools. The simulated radar platform was placed in the upper-left (NW) corner of this dataset, at an altitude of 3000m of the surface of the reference geoid. The result of applying the hidden surface removal algorithm is shown in Figure 3.2, where visible patches are shown in green and hidden patches in black.




Figure 3.1. Grand Canyon Terrain Map




Figure 3.2. Visible Surfaces in the Grand Canyon

One of our objectives is to show how one might leverage previous results in adaptive detection and structured covariance estimation using a geography-based radar simulation tool. Traditional STAP processing of the simulated data has not been performed, because the terrain and land use data lead to a highly nonstationary model for the clutter covariance, when considered as a function of range, and hence the usual paradigm for adaptive detection involving secondary data vectors sharing a common distribution is not really applicable. Furthermore, the usual STAP algorithms use range and Doppler, i.e., the radar coordinates, and our growing familiarity with the geography-based systems, where patches are defined in terms of fixed data points specified in NED datasets, led us to think in ground-based coordinates (see Section 2 above).

We are interested in considering statistical inference problems where the state of nature θ (existence of targets, say) was attached to ground objects such as the patches, independent of the location, orientation, and velocity of the radar platform. The observations and work we have pursued on this contract highlight several potentially fruitful areas of research on such inference problems. Some can be considered generalizations or transformations of standard signal processing algorithms (STAP, GMTI, etc.) Others are inherently new due to the novelty of the problem formulation.

We were also interested in investigating ways that one might exploit the electronic agility of the radar to adaptively *choose* measurements or radar transmit parameters such as waveform or beampattern, so that the radar measurements were maximally informative about the

desired θ . It may very well be that the problem of *experimental design* is just as important as the problem of *data processing* in radar, even though it has received far less attention to date. Accordingly, we began exploring this relatively unknown territory (for us at least) through a literature search and several simple "thought experiments." This eventually led to some very encouraging results in what we now term *active-testing surveillance*. Our initial results in this area are described in [6].

5. Conclusions and Future Work

The results of this study are both encouraging and intellectually stimulating. We now feel comfortable with radar simulation using terrain data, based on our admittedly simple reflectivity model, and feel that this experience could be extended to include more sophisticated EM models, and incorporate more features from geographical information systems, beyond facted-Earth terrain and land-use models. Likewise, we are encouraged by our initial results in active-testing surveillance and feel that this line of investigation could be quite fruitful in the study of systems which adaptively attempt to make the best use of their agility on transmit.

Our plans for the immediate future are to bring together these two research results. We will attempt to demonstrate the feasibility of active testing in the context of an airborne radar employing STAP processing and GIS/GPS/INS side information. As in all research and development programs, this will require a series of incremental improvements on our established results. We see the following topics as being of immediate interest:

- Bringing together the radar simulation work with the active-testing algorithms, to demonstrate how the surveillance system can be placed in ground-based coordinates with accurate geographical side information.
- Incorporation of linear constraints in the illumination pattern (as occurs with radar beamformers) and linear mixing in the sensor measurements (as occurs in the radar receiver) into the active-testing surveillance paradigm.
- Modifications to the methodology that allow for differing target and clutter signatures in the same cell (targets will be moving), and for adaptive processing of interference from unknown sources, such as jamming.
- Investigation of the applicability of emerging Markov Chain Monte Carlo (MCMC) to the more complex scenarios envisioned for active-testing surveillance.

- Implementation of proposed algorithms in high-performance embedded computing systems or commodity graphics engines.

Acknowledgement

This work was supported in part by the U.S. Department of the Air Force through a sub-contract from the Science Applications International Corporation.

References

- [1] J. Ward, "Space-time adaptive processing for airborne radar", MIT Lincoln Laboratory Technical Report 1015, December 1994.
- [2] R. Klemm, *Space-Time Adaptive Processing: Principles and Applications*, IEE Radar, Sonar, Navigation and Avionics Series 9, England, 1998.
- [3] D. Munson, J. O'Brien, and W. Jenkins, "A tomographic formulation of spotlight mode synthetic aperture radar", *Proc. IEEE*, vol. 71, pp. 917-925, Aug. 1983.
- [4] D. Rieken and D. Fuhrmann, "Generalizing MUSIC and MVDR for multiple noncoherent arrays", *IEEE Trans. Signal Processing*, in review.
- [5] D. Rieken and D. Fuhrmann, "Generalizing MUSIC and MVDR for distributed arrays", *Proc. 2002 Asilomar Conference on Signals, Systems, and Computers* (Pacific Grove, CA), November 2002.
- [6] D. Fuhrmann, "Active-testing surveillance systems, or playing Twenty Questions with a radar", *Proc. 11th Workshop on Adaptive Sensor Array Processing (ASAP)*, MIT Lincoln Laboratory, Lexington, MA, March 2002.